

**In the Claims:**

The claims are as follows:

1-25. (Canceled)

26. (Previously presented) A method for maintaining a high availability processing environment, said method comprising:

providing a network having a plurality of clusters, each cluster of the network comprising a plurality of identical servers, each cluster of the network being directly connected to at least one other cluster of the network, wherein each pair of clusters directly connected to each other is characterized by each server in a first cluster of the pair of clusters being directly connected to at least one server in a second cluster of the pair of clusters via a communication link; and

providing a control server adapted to monitor an operational status of said communication link, said operational status of the communication link being that said communication link is operational or non-operational, said control server being directly linked to each server of at least two servers in each cluster via a separate communication channel between the control server and each server of the at least two servers in each cluster, wherein there is no coinciding path segment among the separate communication channels between the control server and each server of the at least two servers in each cluster.

27. (Previously presented) The method of claim 26, said method further comprising:

providing a global dataset that includes an identification of each communication link in the

network, said global dataset being accessible to the control server; and

providing a local dataset specific to each cluster of the plurality of clusters, said local dataset including an identification of each communication link in the network to which the servers of said each cluster is coupled for flow of data out of the cluster, said local dataset being accessible to the servers of said each cluster.

28. (Previously presented) A method for maintaining a high availability processing environment, said method comprising:

providing a network having a plurality of clusters, each cluster of the network comprising a plurality of identical servers, each cluster of the network being directly connected to at least one other cluster of the network, wherein each pair of clusters directly connected to each other is characterized by each server in a first cluster of the pair of clusters being directly connected to at least one server in a second cluster of the pair of clusters via a communication link;

providing a control server adapted to monitor an operational status of said communication link, said operational status of the communication link being that said communication link is operational or non-operational, said control server being directly linked to each server of at least two servers in each cluster via a separate communication channel between the control server and each server of the at least two servers in each cluster, wherein there is no coinciding path segment among the separate communication channels between the control server and each server of the at least two servers in each cluster; and

monitoring an operational status of a first communication link between a first server of the first cluster and a second server of the second cluster, said monitoring being performed by the

control server, said monitoring including sending a query signal to the first server, said query signal requesting the first server to send a response signal to the control server indicating the status of the first communication link, said operational status of the first communication link being that said first communication link is operational or non-operational.

29. (Original) The method of claim 28, wherein the first server is adapted to respond to the query signal by sending a prompt signal over the first communication link to the second server, said prompt signal prompting the second server to send a return signal to the first server over the first communication link, said return signal or absence thereof being indicative of the operational status of the first communication link.

30. (Original) The method of claim 28, wherein the first cluster has a load balancer adapted to distribute data traffic uniformly among the servers comprised by the first cluster, and wherein upon the control server receiving the response signal from the first server such that the response signal indicates that the first communication link is non-operational the method further comprises: notifying the load balancer that the first communication link is non-operational, said notifying being performed by the control server.

31. (Original) The method of claim 28, wherein upon the load balancer being notified that the first communication link is non-operational the method further comprises: failing over the first server with respect to the first communication link, said failing over being performed by the load balancer.

32. (Original) The method of claim 28, wherein upon the control server receiving the response signal from the first server such that the response signal indicates that the first communication link is non-operational the method further comprises: informing a service node in the first cluster that the first communication link is non-operational, said informing being performed by the control server.

33. (Original) The method of claim 32, wherein upon the service node being informed that the first communication link is non-operational the method further comprises: making a determination of a cause of the first communication link being non-operational, said making a determination being performed by the service node.

34. (Original) The method of claim 33, wherein upon the service node making said determination of said cause the method further comprises: facilitating making the first communication link operational, said facilitating being performed by the service node.

35. (Original) The method of claim 28, wherein the first cluster has a load balancer adapted to distribute data traffic uniformly among the servers comprised by the first cluster, and wherein upon the control server not receiving the response signal from the first server within a predetermined period of time after sending the query signal to the first server the method further comprises: making a determination that the first server is non-operational and notifying the load balancer that the first server is non-operational, said making said determination that the first server is non-operational and said notifying being performed by the control server.

36. (Original) The method of claim 35, wherein upon load balancer being notified that the first server is non-operational the method further comprises: failing over the first server, said failing over being performed by the load balancer.

37. (Original) The method of claim 35, wherein upon the control server not receiving the response signal from the first server within a predetermined period of time after sending the query signal to the first server the method further comprises: informing a service node in the first cluster that the first server is non-operational, said informing being performed by the control server.

38. (Original) The method of claim 37, wherein upon the service node being informed that the first server is non-operational the method further comprises: making a determination of a cause of the first server being non-operational, said making said determination of said cause being performed by the service node.

39. (Original) The method of claim 38, wherein upon the service node making said determination of said cause the method further comprises: facilitating making the first server operational, said facilitating being performed by the service node.

40. (Previously presented) The method of claim 26, wherein each cluster of the plurality of clusters has a load balancer that is specific to each cluster and is independent of the load balancer of each other cluster, and wherein the method further comprises distributing data traffic uniformly among the servers comprised by each cluster, said distributing for each cluster being performed by

the load balancer of each cluster.

41. (Original) The method of claim 26, wherein the control server is adapted to receive a message from a first server of the first cluster or from a load balancer of the first cluster, said message indicating that an entity is non-operational, said entity being selected from the group consisting of a server of the first cluster and a communication link between the first server of the first cluster and a second server of the second cluster.

42. (Original) The method of claim 41, wherein upon the control server receiving the message the method further comprises: informing a service node of the first cluster that the entity is non-operational, said informing being performed by the control server.

43. (Original) The method of claim 26, wherein the control server is directly linked to a first server of the first cluster and is not directly linked to a second server of the first cluster, wherein the first server is directly connected to the second server, said method further comprising monitoring the operational status of the second server via direct communication with the first server coupled with direct communication between the first server and the second server, said monitoring being performed by the control server, said operational status of the second server being that said second server is operational or non-operational.

44. (Original) The method of claim 26, wherein at least one cluster of the plurality of clusters does not have a load balancer adapted to distribute data traffic uniformly among the servers comprised

by the first cluster.

45. (Original) The method of claim 26, wherein the plurality of clusters includes a web cluster of web servers, an application cluster of application servers, and a database cluster of database servers, the web cluster being directly connected to the application cluster, the application cluster being directly connected to the database cluster, the web cluster adapted to communicate with the database cluster by way of the application cluster functioning as an intermediary cluster between the web cluster and the database cluster.

46. (Original) The method of claim 45, wherein the web cluster has a load balancer adapted to distribute data traffic uniformly among the web servers comprised by the web cluster, wherein the application cluster has a load balancer adapted to distribute data traffic uniformly among the application servers comprised by the application cluster, and wherein the database cluster has a load balancer adapted to distribute data traffic uniformly among the database servers comprised by the database cluster.

47. (Original) The method of claim 45, wherein the web cluster has a load balancer adapted to distribute data traffic uniformly among the web servers comprised by the web cluster, wherein the application cluster has a load balancer adapted to distribute data traffic uniformly among the application servers comprised by the application cluster, and wherein the database cluster does not have a load balancer adapted to distribute data traffic uniformly among the database servers comprised by the database cluster.

48. (Original) The method of claim 26, wherein the plurality of clusters includes a web cluster of web servers and a database cluster of database servers, the web cluster being directly connected to the database cluster, the web cluster adapted to directly communicate with the database cluster.

49. (Original) The method of claim 48, wherein the web cluster has a load balancer adapted to distribute data traffic uniformly among the web servers comprised by the web cluster, and wherein the database cluster has a load balancer adapted to distribute data traffic uniformly among the database servers comprised by the database cluster.

50. (Original) The method of claim 48, wherein the web cluster has a load balancer adapted to distribute data traffic uniformly among the web servers comprised by the web cluster, and wherein the database cluster does not have a load balancer adapted to distribute data traffic uniformly among the database servers comprised by the database cluster.

51-54. (Canceled)

55. (Previously presented) The method of claim 26, wherein the control server is the only control server adapted to monitor the operational status of said communication link and is the only control server that is directly linked to at the least two servers in each cluster via the communication channel between the control server and the at least two servers in each cluster.

56. (Previously presented) The method of claim 26, wherein the plurality of identical servers in



each cluster consists of all identical servers in each cluster, and wherein the at least two servers in each cluster consists of the plurality of identical servers in each cluster.


57-58. (Canceled)

59. (Previously presented) The method of claim 26, wherein each cluster has an associated load balancer adapted to distribute data traffic uniformly among the servers comprised by each cluster, wherein the control server is linked to each load balancer via a separate link path between the control server and each load balancer, wherein all said separate link paths concurrently exist.

60. (Previously presented) The method of claim 59, wherein the plurality of clusters consist of N clusters such that N is at least 3, wherein the N clusters are denoted as clusters  $C_1, C_2, \dots, C_N$  having the associated load balancers  $L_1, L_2, \dots, L_N$ , respectively, and wherein the N clusters and associated load balancers are serially linked together in a serial order of:

$$L_1 — C_1 — L_2 — C_2 — \dots — L_N — C_N.$$

61. (Previously presented) The method of claim 60, wherein the control server is denoted as C and is linked to the clusters  $C_1, C_2, \dots, C_N$  by links  $L_{C1}, L_{C2}, \dots, L_{CN}$ , respectively, and to the associated load balancers  $L_1, L_2, \dots, L_N$  by links  $L_{L1}, L_{L2}, \dots, L_{LN}$ , respectively, as follows:



wherein each link  $L_{C_n}$  between the control server C and the cluster  $C_n$  represents the separate communication channels between the control server C and each server of the at least two servers in each cluster  $C_n$  for  $n = 1, 2, \dots, N$ .


62. (Previously presented) The method of claim 28, wherein each cluster has an associated load balancer adapted to distribute data traffic uniformly among the servers comprised by each cluster, wherein the control server is linked to each load balancer via a separate link path between the control server and each load balancer, wherein all said separate link paths concurrently exist.

63. (Previously presented) The method of claim 62, wherein the plurality of clusters consist of N clusters such that N is at least 3, wherein the N clusters are denoted as clusters  $C_1, C_2, \dots, C_N$  having the associated load balancers  $L_1, L_2, \dots, L_N$ , respectively, and wherein the N clusters and associated load balancers are serially linked together in a serial order of:

$$L_1 — C_1 — L_2 — C_2 — \dots — L_N — C_N.$$

64. (Previously presented) The method of claim 63, wherein the control server is denoted as C and is linked to the clusters  $C_1, C_2, \dots, C_N$  by links  $L_{C1}, L_{C2}, \dots, L_{CN}$ , respectively, and to the associated

load balancers  $L_1, L_2, \dots, L_N$  by links  $L_{L1}, L_{L2}, \dots, L_{LN}$ , respectively, as follows:



wherein each link  $L_{C_n}$  between the control server  $C$  and the cluster  $C_n$  represents the separate communication channels between the control server  $C$  and each server of the at least two servers in each cluster  $C_n$  for  $n = 1, 2, \dots, N$ .